

On 2 August 2026 the European Union's Artificial Intelligence Act reached what its own text calls its general date of application. Six days earlier, on 27 July, an amending regulation had entered into force that moved the Act's rules for high-risk AI systems—those used in hiring, credit scoring, school admissions and similar decisions—from August 2026 to December 2027, and the rules for AI that is a component of, or is itself, a product already regulated by EU safety law, such as toys and medical devices, from August 2027 to August 2028. In a single summer, two statements were both true: Europe's AI rules now apply, and Europe has delayed its AI rules.

The two statements describe different parts of one law, and a reader who meets only one of them will misread the other. Sorting them out requires three ideas that matter well beyond Europe: that a law has several dates, not one; that a binding obligation is a different thing from a voluntary framework, even when the two are written to fit together; and that the reach of a rule is set by where a product is sold, not only by where its maker sits. Together they answer the practical question behind most news about AI regulation—does this rule bind this product, here, today?

Two readings that mislead

The first misreading is that Europe has put its AI law on hold. Much of it was untouched. Prohibitions on certain practices—among them the social scoring of individuals—have applied since 2 February 2025. Obligations on providers of general-purpose AI models, the large models behind chatbots and coding assistants, have applied since 2 August 2025, including the duty to publish a summary of the content used for training. From 2 August 2026 the Commission's AI Office gained its enforcement powers over those providers: the Commission describes them as the power to request technical documentation, evaluate models, require corrective measures and issue fines. The duty on providers of systems that generate synthetic audio, images, video or text to mark their output, in a machine-readable format, as artificially generated also applies from that date, although systems already on the market before it were given until 2 December 2026.

The amending regulation did not only defer. It added a prohibition. From 2 December 2026 the Act bans placing on the market or putting into service AI systems that generate or manipulate realistic sexual images, video or audio of identifiable people without their explicit consent, or child sexual abuse material, where that is the system's intended purpose or a reasonably foreseeable outcome against which it lacks reasonable safeguards; using such a system is banned where the user's purpose is to make that material. The Commission's own summary counts this as the Act's ninth prohibition, "introduced as a part of the AI Omnibus".

The second misreading runs the other way: that the whole AI Act now applies and companies are already being penalised under it. It has been in force since 1 August 2024. Being in force and applying are distinct legal states, and the Act's obligations arrive in tranches that run to 2 August 2028 for AI embedded in

regulated products, and to 2 August 2030 for high-risk systems intended for use by public authorities that were already on the market. A rule that applies, moreover, is not yet a rule under which anyone has been penalised. According to AFP, the Commission said on 1 September 2026 that it had sent requests for information to more than 30 companies in the AI sector, which the agency described as "a preliminary step before the possible launch of a formal investigation". No fine under the Act had been reported in the sources consulted.

Headlines fed both readings. On 19 November 2025, reporting the Commission's proposal, the Greek daily *To Vima* ran "EU Delays 'High Risk' AI Rules to 2027 After Tech Pushback"—a proposal described as though it were law. On 4 August 2026 *Help Net Security* headlined "EU begins enforcing AI Act, putting AI models under the microscope", accurately echoing the Commission's own release of 31 July, "Commission starts enforcing AI Act rules and new transparency requirements on 2 August"; in its first month on the record, that enforcement amounted to a round of information requests, not penalties.


A law has more than one date

A European regulation typically passes through four dated stages. It is adopted and signed; it is published in the Official Journal; it enters into force, usually on the twentieth day after publication; and it applies, either at once or from a later date the text sets. Entry into force makes the act part of the law and starts its clocks running. Application is when its duties bind the people they address.

The General Data Protection Regulation is the familiar example. It entered into force in May 2016 and, in the words of its Article 99, "shall apply from 25 May 2018"—two years in which firms were expected to prepare. The AI Act's final article follows the same pattern and then complicates it:

This Regulation shall enter into force on the twentieth day following that of its publication in the Official Journal of the European Union. It shall apply from 2 August 2026. However:

The "however" introduces the tranches. As signed in June 2024, Article 113 brought the general provisions and prohibitions forward to 2 February 2025; the governance chapter, the rules for general-purpose models and the penalties chapter to 2 August 2025, with the exception of the Commission's power to fine model providers; and the rules for AI embedded in regulated products back to 2 August 2027. The amending regulation of July 2026 rewrote two of those points and added a third.

Figure 1. Months from the AI Act's entry into force (1 August 2024) to the date each set of rules applies

Sources: Regulation (EU) 2024/1689, Article 113 (Official Journal L, 12 July 2024); Regulation (EU) 2026/1744, Article 1, point 40, amending Article 113 (Official Journal L, 24 July 2026). Blue bars are dates as adopted in 2024 that the 2026 amendment left in place or replaced; orange bars are dates set by the 2026 amendment. Months are counted from 1 August 2024 and rounded to the nearest whole month.

▼ Table view

Figure 1. Months from the AI Act's entry into force (1 August 2024) to the date each set of rules applies

Item	Value
Prohibitions and AI literacy (2 Feb 2025)	6 months
General-purpose model obligations (2 Aug 2025)	12 months
General date of application; Commission fines for model providers (2 Aug 2026)	24 months
High-risk uses listed in Annex III, as adopted (2 Aug 2026)	24 months
High-risk uses listed in Annex III, as amended (2 Dec 2027)	40 months
AI in or as regulated products (Annex I), as adopted (2 Aug 2027)	36 months
AI in or as regulated products (Annex I), as amended (2 Aug 2028)	48 months
New prohibition added in 2026 (2 Dec 2026)	28 months

Two further dates matter for products already on the market. Providers of general-purpose models placed on the market before 2 August 2025 have until 2 August 2027 to comply. High-risk systems already in service before their rules apply are caught only if they are significantly changed afterwards—except those intended for use by public authorities, which must comply by 2 August 2030 in any case.

The bargain of 1985

The dates moved because of a design choice made long before anyone was writing law about AI. In February 1979 the Court of Justice's *Cassis de Dijon* judgment (Case 120/78) held, as the Commission's Blue Guide reads it, that a member state could keep out a product lawfully sold elsewhere in the Community only on the

ground that it failed an essential requirement. That forced a question: what, then, should Community law on products actually contain? The answer was the New Approach, a technique "approved by the Council of Ministers on 7 May 1985", in the words of the Commission's guide to EU product rules, the "Blue Guide". The Council's resolution set out four principles. Legislation would be limited to the essential requirements products must meet; the technical specifications would be entrusted to standards organisations; and, in the resolution's own words,


these technical specifications are not mandatory and maintain their status of voluntary standards

while national authorities would be obliged to presume that products built to them conform. A producer could decline to follow a standard, the resolution added, but "in this event he has an obligation to prove that his products conform to the essential requirements". The Blue Guide's summary, written in 2022, is that the New Approach

restricted the content of legislation to 'essential requirements' leaving the technical details to European harmonised standards

The law states what a product must achieve. European standardisation bodies, at the Commission's request, write the technical specifications for achieving it. When the Commission publishes the reference of such a harmonised standard in the Official Journal, a product built to it is presumed to conform to the corresponding requirements. The guide is careful about what is and is not compulsory: "The application of harmonised or other standards remains voluntary", and a manufacturer may meet the requirements another way, but then carries the burden of demonstrating that it has. The New Legislative Framework of July 2008 built on the same design and added rules for conformity assessment, accreditation and market surveillance.

The AI Act is written in that tradition. Its recitals say its rules for high-risk systems are laid down "consistently with" the New Legislative Framework, and its Article 40 provides that systems in conformity with published harmonised standards "shall be presumed to be in conformity" with the Act's requirements for high-risk systems, to the extent the standards cover them. That design concentrates risk in one place. Where standards are not ready when obligations arrive, providers must show compliance by other means—common specifications, if the Commission has adopted them under Article 41, or their own technical solutions—and national authorities must judge those showings with less of a common yardstick.

Figure 2. How a European product rule becomes something a firm can comply with

Schematic of the New Approach (1985) and New Legislative Framework (2008) as applied in Regulation (EU) 2024/1689, Articles 40, 41 and 53(4). Sources: Commission notice 2022/C 247/01 (the 'Blue Guide'), section 1.1; Regulation (EU) 2024/1689; Regulation (EU) 2026/1744, recital 41.

▼ Table view

Figure 2. How a European product rule becomes something a firm can comply with — stages

#	Stage	Note
1	Legislature	Parliament and Council adopt the essential requirements
2	Commission request	asks the standards bodies for technical specifications
3	Harmonised standard	written by European standardisation bodies; use is voluntary
4	Reference published in the Official Journal	
5	Presumption of conformity	a product built to the standard is presumed to comply
6	Lawful sale across the single market	
7	Common specifications	Commission implementing act, Article 41; presumption to the extent covered
8	Code of practice	general-purpose models, Article 53(4); no presumption (2026/1744, recital 41)

Figure 2. How a European product rule becomes something a firm can comply with — connections

From	To	Label
Legislature	Commission request	
Commission request	Harmonised standard	
Harmonised standard	Reference published in the Official Journal	
Reference published in the Official Journal	Presumption of conformity	
Presumption of conformity	Lawful sale across the single market	
Legislature	Common specifications	if standards are late
Common specifications	Presumption of conformity	to the extent covered
Legislature	Code of practice	general-purpose models
Code of practice	Lawful sale across the single market	may be relied on to show compliance

That is what happened. The amending regulation gives its reason in its own recital 40: for the high-risk obligations,

the delayed availability of standards, common specifications, and alternative guidance and the delayed establishment of national competent authorities lead to challenges that jeopardise the effective entry into application of those obligations

—challenges, the recital continues, that "risk a significant increase in implementation costs in a way that does not justify maintaining their initial date of application, namely 2 August 2026". The stated causes are two—missing standards and their substitutes, and authorities not yet in place—and the stated harm is cost.

The remedy changed on the way through. The Commission's proposal of 19 November 2025 would have made the high-risk rules apply six or twelve months after "a decision of the Commission confirming that adequate measures in support of compliance with Chapter III are available", or on fallback dates of 2 December 2027 and 2 August 2028 if those came first. When the European Parliament adopted its position on 26 March 2026, by 569 votes to 45, its press release said that "MEPs introduce fixed dates for application


to ensure predictability and legal certainty". The final regulation contains only the fixed dates. The high-risk rules do not, as enacted, wait for the standards; the standards' lateness was the reason the fixed dates were set later.

Obligation, framework, and the space between

An obligation is a duty imposed by law, with an authority empowered to enforce it and a penalty attached. A framework is a document in which a company or a group of companies describes how it intends to behave. The first can be enforced against a firm that disagrees with it; the second can be revised by its author.

The AI Act's penalties show what the obligation side looks like. Fines are capped as the higher of a fixed sum and a share of worldwide annual turnover, and the share depends on the infringement.

Figure 3. Maximum fines under the AI Act, as a share of worldwide annual turnover



Source: Regulation (EU) 2024/1689, Articles 99 and 101. Each cap is the higher of the percentage and the fixed sum for undertakings. For SMEs the lower of the two applies, and the 2026 amendment extends that treatment to small mid-cap enterprises for the 3% and 1% tiers of Article 99. Article 101 fines require the Commission to find that the provider acted intentionally or negligently. These are ceilings, not typical fines.

▼ Table view

Figure 3. Maximum fines under the AI Act, as a share of worldwide annual turnover

Item	Value
Prohibited practices (or EUR 35m if higher), Article 99(3)	7%
Most other operator obligations, including transparency (or EUR 15m), Article 99(4)	3%
General-purpose model providers, fined by the Commission (or EUR 15m), Article 101	3%
Incorrect or misleading information to authorities (or EUR 7.5m), Article 99(5)	1%

Between the two sits a hybrid that the Act creates on purpose. Article 56 asks the AI Office to facilitate codes of practice, and Article 53(4) lets providers of general-purpose models "rely on codes of practice ... to demonstrate compliance" with their obligations "until a harmonised standard is published". The General-Purpose AI Code of Practice, published on 10 July 2025, is described by the Commission as

a voluntary tool, prepared by independent experts in a multi-stakeholder process, designed to help industry comply with the AI Act's obligations for providers of general-purpose AI models

Signing it is voluntary; the obligations it helps a provider meet are not. The amending regulation is explicit about the code's weight: such codes "have limited legal effect, and in particular do not grant a presumption of conformity", which is why it removed the Commission's power to approve them by implementing act. The Commission's list of signatories, read on 10 October 2026, includes Amazon, Anthropic, Google, IBM, Microsoft, Mistral AI and OpenAI, among others; xAI signed only the code's safety and security chapter. Meta does not appear on the list. A provider outside the code must show compliance by other means—the same choice a manufacturer faces when it declines to follow a harmonised standard.

At the far end of the spectrum sit the companies' own frontier-safety frameworks, such as Anthropic's Responsible Scaling Policy, first published in September 2023, and comparable documents at OpenAI and Google DeepMind. These are promises a company makes about itself, and its author can change them. Anthropic's own page, read on 10 October 2026, lists five versions of its policy taking effect between February and July 2026, the latest on 8 July; Google DeepMind published the third iteration of its Frontier Safety Framework in September 2025 and updated it on 17 April 2026. Anthropic makes Claude, the model with which the podcast and its written edition are produced, and each company's page is the only source for its own revisions.

Laws are revised too; what differs is who revises them and by what procedure. The July 2026 regulation rewrote an obligation that had applied since February 2025. Article 4 had required providers and deployers to ensure, "to their best extent, a sufficient level of AI literacy" among their staff; it now requires them "to take measures to support the development of AI literacy", and adds that the obligation "does not require providers or deployers to guarantee any specific level of AI literacy of any individual". That change took a proposal, a parliamentary vote and a Council decision. A company's framework changes when the company decides.

The companies' frameworks began as soft commitments to governments. On 21 July 2023 the White House announced "voluntary commitments" from seven companies. At the AI Seoul Summit in May 2024, a larger group—including Amazon, Anthropic, Google, Meta, Microsoft, OpenAI and xAI—undertook to act "in accordance with the following voluntary commitments, and to demonstrate how they have achieved this by publishing a safety framework focused on severe risks".

California then moved the same kind of document across the line. Senate Bill 53, the Transparency in Frontier Artificial Intelligence Act, approved on 29 September 2025 as Chapter 138 of the Statutes of 2025, adds to the state's Business and Professions Code a section providing that a large frontier developer


shall write, implement, comply with, and clearly and conspicuously publish on its internet website a frontier AI framework

describing how it approaches matters including the incorporation of national and international standards and industry-consensus best practice. A developer may change its framework, but must publish any material modification "and a justification for that modification within 30 days". A developer that "fails to comply with its own frontier AI framework"—or fails to publish a required document, makes a materially false or misleading statement about catastrophic risk or about its compliance with its framework, or fails to

report a critical safety incident—faces a civil penalty of up to \$1m per violation, "recovered in a civil action brought only by the Attorney General". The state supplies the duty and the enforcer; the developer still supplies most of the content.

New York followed. Governor Kathy Hochul signed the RAISE Act on 19 December 2025, and a chapter amendment signed on 27 March 2026 rewrote its central duty in the same words as California's—a large frontier developer must write, follow and publish a frontier AI framework—and moved its effective date to 1 January 2027.

Figure 4. From promise to obligation: four kinds of AI rule



Schematic. The examples are dated in the text: the EU code of practice, 10 July 2025; SB 53, 29 September 2025; Regulation (EU) 2024/1689, 13 June 2024.

▼ Table view

Figure 4. From promise to obligation: four kinds of AI rule — stages

#	Stage	Note
1	Company framework	written and revised by its author; e.g. lab safety policies
2	Voluntary code with limited legal effect	EU general-purpose AI code: signing optional; adherence shows compliance
3	Statute requiring a framework	California SB 53: write, implement, comply with and publish; Attorney General enforces
4	Statute setting the duties	EU AI Act: requirements in law; fines up to 7% of turnover

Figure 4. From promise to obligation: four kinds of AI rule — connections

From	To	Label
Company framework	Voluntary code with limited legal effect	more binding
Voluntary code with limited legal effect	Statute requiring a framework	more binding
Statute requiring a framework	Statute setting the duties	more binding

Whose rules reach a product

Jurisdiction answers a prior question: whose rules apply at all. The AI Act attaches to the market rather than to the maker. Article 2(1)(a) applies it to providers placing AI systems or general-purpose models on the market in the Union,

irrespective of whether those providers are established or located within the Union or in a third country

and Article 2(1)(c) extends it to providers and deployers outside the Union "where the output produced by the AI system is used in the Union". A model developer in California that offers its product in Europe is bound by Brussels for that product, as a car maker in Japan is bound by European type-approval rules for the cars it sells in Germany.

The United States has no federal statute of this kind. Rules have come from the states—California's SB 53 among them—and the federal government has set out to challenge them. Executive Order 14365, signed on 11 December 2025, directs the Attorney General to establish an AI Litigation Task Force

whose sole responsibility shall be to challenge State AI laws inconsistent with the policy set forth in section 2 of this order

on grounds including that such laws "unconstitutionally regulate interstate commerce" or are preempted by existing federal regulations. It instructs the Commerce Department to identify "onerous" state AI laws and to make states that keep them ineligible, to the extent federal law allows, for certain broadband funds. And it asks two presidential advisers to prepare "a legislative recommendation establishing a uniform Federal policy framework for AI that preempts State AI laws that conflict with the policy set forth in this order"—with carve-outs, among them child-safety laws. The order also directs the Federal Communications Commission to consider a federal reporting and disclosure standard "that preempts conflicting State laws", and the Federal Trade Commission to explain when state laws requiring altered outputs are preempted. The order can direct federal lawyers, federal money and federal agencies; it cannot by itself displace a state law—that takes an act of Congress or a valid federal rule that preempts it, or a court ruling that it is invalid. On 20 March 2026 the White House published what it called a "National AI Legislative Framework". No federal statute preempting state AI laws had been enacted by 10 October 2026 in the record consulted; a House bill to that effect, H.R. 5388, introduced in September 2025, had not moved past introduction.

The first courtroom test came in Colorado. On 9 April 2026 xAI sued the state's attorney general over Colorado's 2024 AI Act, SB 24-205, whose duties had originally been due to start on 1 February 2026 and which a law approved on 28 August 2025, SB 25B-004, had already deferred to 30 June 2026. On 24 April the United States intervened, unopposed, in what the court described as an equal-protection case, the Acting Attorney General having certified it to be of general public importance. On 27 April the court granted a joint motion under which the state "shall not initiate enforcement" of the law, or of any replacement enacted that session, for violations occurring up to 14 days after the court rules on a preliminary injunction. On 14 May Colorado's governor signed SB 26-189, which "repeals and reenacts" the 2024 provisions with new requirements on automated decision-making technology, starting on 1 January 2027. No court has ruled on the merits.

From signature to application, step by step

Date	Event	What it changed
13 June 2024	AI Act signed (Regulation (EU) 2024/1689)	Published 12 July 2024; in force 1 August 2024
2 February 2025	First tranche applies	Prohibitions; AI literacy
10 July 2025	General-Purpose AI Code of Practice published	Voluntary route to show compliance
2 August 2025	Second tranche applies	General-purpose model obligations; governance; penalties chapter
29 September 2025	California SB 53 signed	Large frontier developers must write, implement, comply with and publish a framework
19 November 2025	Commission proposes the "AI Omnibus"	Proposal only
11 December 2025	US Executive Order 14365	Litigation task force against state AI laws
26 March 2026	European Parliament adopts its position, 569 to 45	Proposes fixed dates in place of the conditional trigger
7 May 2026	Political agreement on the Omnibus	Not yet law
16 and 29 June 2026	Parliament's first-reading position; Council's decision	Adoption
8 July 2026	Regulation (EU) 2026/1744 signed	Not yet in force
24 July 2026	Published in the Official Journal	In force on the third day after
27 July 2026	Amendment enters into force	High-risk dates moved; new prohibition dated 2 December 2026
2 August 2026	General date of application	AI Office enforcement powers over model providers begin
1 September 2026	Commission announces its first requests for information, per AFP	More than 30 companies; a preliminary step, not a fine

The amending regulation's recital 46 explains the haste: it was to enter into force "as a matter of urgency on the third day following that of its publication", in order to ensure legal certainty "with a view to the imminent general application" of the Act. An amendment published on 24 July that had followed the ordinary twenty-day rule would have entered into force after the date it was written to move.

Mapped onto the ideas above, the pattern largely holds. The rules that kept their dates had a finished instrument for compliance: providers of general-purpose models could rely on the code of practice until standards existed. The high-risk rules that moved were those whose standards and planned substitutes, by the amendment's own account, were late, and whose national authorities were not all in place. Read together, the documents suggest the pattern rather than state it. The marking duty for AI-generated content also had a code of its own (the Commission's opinion on it is dated 9 July 2026) and kept its August 2026 date for new systems, though systems already on the market were given four more months, to 2 December

2026, so that providers could adapt their practices "within a reasonable time without disrupting the market". And the Commission says the new high-risk dates mean "the rules apply when companies have the right support tools to facilitate implementation, such as standards".

Two postures

The European posture is that of one lawmaking process—the Parliament and the member governments' Council—binding a market of 27 countries through a single regulation, then amending its own timetable—by the same ordinary legislative procedure—when the machinery it depends on is not ready.

The American posture, as of October 2026, is fragmented. States legislate; the federal executive challenges them in court and with funding conditions, and asks Congress for a single national framework that would displace conflicting state rules. Until Congress acts or a court rules, a company selling in California is bound by California's statute whatever federal policy prefers.

The two postures can meet inside one firm. A developer that meets California's thresholds for a large frontier developer must, by statute, write, follow and publish a frontier AI framework; in Europe the duties for the most advanced models are statutory too, and following the code's safety and security chapter is a voluntary way to show they are met.

What to watch

2 December 2026. The new prohibition on non-consensual sexual imagery of identifiable people and on child sexual abuse material applies, and generative systems already on the market before August must mark their outputs as artificially generated. The first public enforcement step under either—by the Commission or a national authority—would show whether "applies" has become "enforced".

1 January 2027. New York's RAISE Act takes effect, and the duties in Colorado's replacement law begin, with the Colorado attorney general's rules on post-adverse-outcome disclosures due by the same date. Under the April court order, xAI's request for a preliminary injunction is due within 28 days of those rules being finally adopted.

2 August 2027. General-purpose models placed on the market before 2 August 2025 must comply with the EU's model obligations.

Before 2 December 2027. Whether references to harmonised standards for high-risk AI systems are published in the Official Journal. If they are, a provider building to them is presumed to conform when the Annex III rules apply; if they are not, the question is whether the date moves a second time.

The idea to keep

A rule binds a product only when three things line up: a duty rather than a promise, a reach that covers the market where the product is sold, and a date that has actually arrived—with "in force", "applies" and "enforced" kept distinct. Europe's summer of 2026 showed all three at work at once. The duties for general-purpose models, for which the code of practice offered a finished compliance route, kept their timetable; the duties for high-risk systems moved because, on the Union's own account, the standards and authorities they relied on were not ready. The design behind both is older than any chatbot: the 1985 bargain in which law sets essential requirements and standards supply the detail. Asked of any new AI rule, the three questions are who stands behind it—a law with an enforcer, or only a company's word—whose market it reaches, and which of its dates has arrived.

Sources

Source	Date
Regulation (EU) 2016/679 (General Data Protection Regulation), Article 99	Official Journal 4 May 2016
European Commission, <i>The "Blue Guide" on the implementation of EU product rules 2022</i> , Commission notice 2022/C 247/01, Official Journal C 247, section 1.1	29 June 2022
Regulation (EU) 2024/1689 (Artificial Intelligence Act), recital 9; Articles 2, 4, 5, 6, 40, 41, 50, 53, 56, 99, 101, 111 and 113	signed 13 June 2024; Official Journal 12 July 2024
European Commission, <i>The General-Purpose AI Code of Practice</i> (page and signatory list)	code published 10 July 2025; read 10 October 2026
California Senate Bill 53, <i>Transparency in Frontier Artificial Intelligence Act</i> , Chapter 138, Statutes of 2025, sections 22757.12 and 22757.15 of the Business and Professions Code as added	approved 29 September 2025
Office of the Governor of California, <i>Governor Newsom signs SB 53</i>	29 September 2025
Executive Order 14365, <i>Ensuring a National Policy Framework for Artificial Intelligence</i> , 90 FR 58499, sections 3 to 8	signed 11 December 2025; published 16 December 2025
Regulation (EU) 2026/1744 (Digital Omnibus on AI), recitals 38, 40, 41 and 46, Articles 1 and 4	signed 8 July 2026; Official Journal 24 July 2026
European Commission, <i>Regulatory framework for AI</i> (AI Act policy page)	read 10 October 2026
Anthropic, <i>Responsible Scaling Policy</i> page (version history)	last updated 14 August 2026; read 10 October 2026
Council Resolution of 7 May 1985 on a new approach to technical harmonization and standards (85/C 136/01), Official Journal C 136, Annex II	7 May 1985; published 4 June 1985
European Commission, proposal COM(2025) 836 (Digital Omnibus on AI), draft Article 113(d)	19 November 2025
European Parliament, press release 20260323IPR38829, <i>Artificial Intelligence Act: delayed application, ban on nudifier apps</i>	26 March 2026
White House, fact sheet on voluntary AI commitments	21 July 2023
UK Government, <i>Frontier AI Safety Commitments, AI Seoul Summit 2024</i>	May 2024
Office of the Governor of New York, RAISE Act signing release; New York Senate bill S8828 (Chapter 96 of 2026), section 3	19 December 2025; 27 March 2026

Source	Date
<i>X.AI LLC v. Weiser</i> , No. 1:26-cv-01515 (D. Colo.), docket and orders of 24 and 27 April 2026	filed 9 April 2026; read 10 October 2026
Colorado General Assembly, SB 26-189 <i>Automated Decision-Making Technology</i> , bill page	signed 14 May 2026
Google DeepMind, <i>Strengthening our Frontier Safety Framework</i>	22 September 2025, updated 17 April 2026
AFP, <i>EU questions dozens of companies using new AI powers</i> (via The Star)	2 September 2026
To Vima, <i>EU Delays ‘High Risk’ AI Rules to 2027 After Tech Pushback</i> ; Help Net Security, <i>EU begins enforcing AI Act, putting AI models under the microscope</i>	19 November 2025; 4 August 2026
GovTrack, H.R. 5388 status page	read 10 October 2026
European Commission, press release <i>Commission starts enforcing AI Act rules and new transparency requirements on 2 August</i>	31 July 2026
The White House, <i>President Donald J. Trump Unveils National AI Legislative Framework</i>	20 March 2026
Colorado General Assembly, SB 25B-004, bill page	approved 28 August 2025